Cost Impact Analysis

A Regression-Based Business Optimisation Methodology for Connecting Organisational Expenditure to Measurable Outcomes

DOI: 10.5281/zenodo.20713675

Daniel Diosi
Daniel Diosi & Partners
ORCID: 0009-0003-1051-3117
costimpactmonitor.com

Abstract 1
0. Capital Misallocation as an Economic Problem 3
0.1 The Scale of the Problem 3
0.2 The Productive Investment Displacement Effect 4
0.3 The Platform Concentration Effect 4
0.4 The SMB Productivity Gap 5
0.5 The Public Sector Efficiency Argument 5
1. The Contemporary Business Information Problem 7
1.1 The Cost Inflation Trend 7
1.2 Reporting Fragmentation 7
1.3 Information Overload and Underutilisation 8
1.4 The Missing Function 8
2. Evidence Across Sectors 10
2.1 Private Sector 10
2.2 Public Sector 10
2.3 Civil Society 11
3. Critique of Existing Approaches 12
3.1 Platform and Channel Reporting 12
3.2 Business Intelligence Infrastructure 12
3.3 Marketing Mix Modelling 13
3.4 Management Accounting Frameworks 13
3.5 The Complexity Trap 14
4. The Proposed Methodology — Cost Impact Analysis 15
4.1 The Business Optimisation Context 15
4.2 The Core Proposition 15
4.3 Analytical Architecture — A Two-Layer Statistical Framework 15
4.3.1 Primary Layer — Multivariate Ordinary Least Squares Regression 16
4.3.2 Secondary Layer — Cost Flow Correlation on Transformed Series 18
4.4 Application Requirements 20
4.5 Constraints and Limitations 20
4.6 The Spend Efficiency Construct 21
5. Implications for Business Efficiency & Optimisation 22
5.1 The Executive Intelligence Layer 22
5.2 Sector-Specific Implications 22
5.3 The Democratisation of Analytical Capability 22
5.4 Contribution to Business Optimisation Literature 23
Conclusion 25
References 26

Abstract

The inability of organisations to connect expenditure to measurable outcomes constitutes a documented market inefficiency with consequences extending beyond individual management failure to aggregate capital misallocation, platform concentration, and a structural productivity gap between analytically capable and analytically constrained organisations. Despite the proliferation of business intelligence infrastructure, platform reporting, and management accounting frameworks, no accessible, cross-departmental methodology exists for systematically connecting categorised organisational costs to defined business outcome KPIs — independent of the departments and vendors whose performance is being assessed.

This paper proposes Cost Impact Analysis as a named business optimisation discipline that addresses this gap through a deliberate two-layer statistical architecture. The primary layer applies multivariate ordinary least squares regression to raw aligned period-level cost and KPI data, estimating the independent contribution of each cost category to outcome variation while controlling for co-movement between categories. The secondary layer applies Pearson correlation to detrended, first-differenced series, providing a significance-filtered visual representation of co-movement strength after trend removal. A composite Spend Efficiency Score synthesises multiple dimensions of cost-outcome performance into a single directional executive indicator.

The methodology is designed for deployment across private, public, and civil society organisations without dedicated analytical infrastructure, addressing a structural accessibility gap that existing approaches — including Marketing Mix Modelling, enterprise BI platforms, and management accounting frameworks — do not resolve. The paper situates Cost Impact Analysis within the business optimisation literature and argues that its widespread adoption would constitute a measurable contribution to economic capital allocation efficiency.

0. Capital Misallocation as an Economic Problem

The inability of organisations to connect expenditure to outcomes is not merely an internal management failure. At aggregate scale it constitutes a documented market inefficiency with measurable macroeconomic consequences — one that directs capital toward the most self-reported activities rather than the most productive ones, systematically transfers economic value from productive business investment to platform intermediaries, and widens the productivity gap between organisations with sophisticated analytical capability and those without it.

This paper proposes Cost Impact Analysis as a business optimisation methodology that addresses this inefficiency — not only as a tool for individual organisational improvement but as a contribution to the broader project of making capital allocation decisions more evidence-based across the economy. The argument that follows is therefore simultaneously a management methodology proposal and an economic efficiency argument.

0.1 The Scale of the Problem

Digital advertising costs have risen consistently and significantly above general price inflation. Google cost-per-click increased by approximately 18 percent year-over-year between 2021 and 2023, with the annual rate moderating to approximately 14 percent by 2024 while reaching record highs in e-commerce, insurance, and technology sectors (TruLata Solutions, 2025). Meta CPMs rose 20.03 percent year-over-year from 2024 to 2025, with every single industry tracked experiencing an increase — ranging from 8.08 percent in the most contained categories to 38.03 percent in Health and Wellness (Triple Whale, 2025). These increases consistently and materially outpace general consumer price inflation, which averaged approximately 4 percent annually in major economies over the same period.

The SaaS subscription layer has grown in parallel. The average company now operates more than 130 SaaS applications, with 44 percent of SaaS licenses going unused or underutilised — representing an estimated 18 billion US dollars in annual waste across the market (Colorlib, 2026; BetterCloud, 2024). For SMBs specifically, companies with 75 to 199 employees used an average of 44 SaaS applications in 2024, with nearly 50 percent of licenses going unused for 90 days or more (SellersCommerce, 2026). The average company wastes an estimated 135,000 US dollars annually in unused SaaS licenses alone (Digital Silk, 2025).

Agency retainers add a further compounding cost layer — typically increasing 5 to 10 percent annually through a combination of explicit rate increases and scope expansion — while remaining structurally excluded from the ROI calculations presented to business leadership.

0.2 The Productive Investment Displacement Effect

Capital spent on digital advertising at inflating platform costs, unverified agency retainers, and unused technology subscriptions is capital not invested in product development, workforce capability, or operational improvement. When this misallocation occurs simultaneously across thousands of organisations it represents a measurable drag on productive investment at the macroeconomic level.

The economic literature on capital misallocation documents this phenomenon in multiple contexts. Hsieh and Klenow (2009) demonstrated that misallocation of capital and labour across firms in China and India explained a significant portion of the total factor productivity gap with the United States. Syverson (2011) identified measurement capability and information availability as among the primary determinants of productivity differences across firms. The extension of this insight to the specific context of cost-outcome measurement failure in digital and operational expenditure is the contribution this paper advances.

When organisations cannot distinguish which expenditure produces results, capital allocation follows the path of least resistance — toward activities that generate *visible metrics* rather than *measurable outcomes*. Platform advertising generates reports. Agency retainers produce presentations. Tool subscriptions produce dashboards. Each creates the appearance of productive activity while the question of whether that activity produced a proportionate business result remains systematically unanswered.

0.3 The Platform Concentration Effect

The systematic inability of businesses to verify the returns on platform advertising creates a structural transfer of capital from productive business investment to platform intermediaries — a transfer that accelerates as platforms raise prices knowing that their customers lack the measurement infrastructure to evaluate the return on that expenditure.

This concentration of economic value in platform intermediaries has been documented in the competition economics literature. Philippon (2019) identified the role of reduced competition in enabling sustained price extraction by dominant platform businesses. The measurement failure that enables the cost inflation documented in Section 0.1 is a structural condition that reinforces this concentration — organisations that cannot measure the return

on platform spend cannot make rational switching or reduction decisions, sustaining demand regardless of actual value delivered.

The aggregate consequence is a systematic and sustained transfer of economic value from productive business operations to platform intermediaries — representing not a market failure in the traditional sense but a measurement failure with market-distorting consequences at scale.

0.4 The SMB Productivity Gap

Small and medium-sized businesses represent the majority of employment and a significant share of economic output in most developed economies. The OECD has documented a persistent and growing productivity gap between large enterprises and SMBs across member countries (OECD, 2021), attributing part of this gap to the differential availability of management practices, analytical capability, and decision-making infrastructure between large and small organisations.

Cost-outcome measurement capability is a specific instance of this broader analytical gap. Large enterprises with dedicated data functions, business intelligence infrastructure, and analytical resource can implement sophisticated cost-outcome measurement frameworks. SMBs without this resource cannot. The consequence is that SMBs systematically allocate capital with less evidence than their larger competitors — not because their decision-makers are less capable, but because the measurement tools available to them are less adequate.

A methodology that makes rigorous cost-outcome analysis accessible to SMBs — within the resource and capability constraints that characterise these organisations — addresses not only a business management gap but a structural productivity disadvantage with economic development implications. The democratisation of analytical capability is, at aggregate scale, a contribution to economic productivity.

0.5 The Public Sector Efficiency Argument

In public sector contexts the economic stakes of measurement failure are directly quantifiable in terms of public resource. Capital allocated to programmes, communications functions, and operational activities without systematic outcome measurement represents a transfer of public funds with unverifiable public benefit — a condition that audit frameworks increasingly identify but analytical infrastructure rarely resolves.

At the scale of national and regional government expenditure the aggregate cost of this measurement gap is material. The growing pressure from central government bodies, audit

institutions, and public accountability frameworks for value-for-money evidence exists in an environment where the analytical frameworks available to most public sector finance functions are oriented toward expenditure compliance rather than outcome verification. Cost Impact Analysis addresses this gap with a methodology calibrated for the resource constraints and accountability requirements of public sector organisations.

The economic argument established in this section frames the remainder of the paper. The cost-outcome measurement gap described in Sections 1 through 3 is not a collection of isolated organisational failures — it is a systemic market condition with documented macroeconomic consequences. The methodology proposed in Section 4 is a response not only to an organisational need but to an economic efficiency problem.

1. The Contemporary Business Information Problem

Organisations are experiencing a paradox — more data, more reporting, and more analytical infrastructure than at any previous point in commercial history, while simultaneously lacking a systematic, independent framework for answering the most fundamental question in financial management: did this expenditure produce a result? This section documents the specific manifestations of that paradox that motivate the methodology proposed in Section 4.

1.1 The Cost Inflation Trend

The expenditure environment within which organisations operate has changed materially over the past decade. The costs associated with reaching customers, developing people, and operating infrastructure have increased — in some categories substantially — while the measurement frameworks available to evaluate the return on those costs have not kept pace.

Digital advertising cost inflation is among the most documented and most consequential of these trends. Google cost-per-click increased approximately 10 percent year-over-year in 2024 across 86 percent of industries (WordStream, 2024; PPC Hero, 2025). Over a three-year period from 2021 to 2024, the compound increase reached 18 percent in multiple sectors, with record highs recorded in e-commerce, insurance, and technology. Meta CPMs rose 20.03 percent year-over-year from 2024 to 2025, with every single industry tracked experiencing an increase (Triple Whale, 2025). These increases consistently outpace general consumer price inflation, representing a real reduction in advertising value for equivalent nominal expenditure.

The SaaS subscription layer compounds this trend. The average organisation now operates more than 130 SaaS applications (Colorlib, 2026), with 44 percent of licenses going unused or underutilised — representing an estimated 18 billion US dollars in annual waste. Agency retainers increase annually through a combination of explicit rate increases and quiet scope expansion. The aggregate consequence is a cost base that grows consistently faster than the measurement capability available to evaluate it.

1.2 Reporting Fragmentation

The costs described above are not reported in a unified framework. They are distributed across multiple systems, departments, payment methods, and reporting cycles — each producing its own view of performance without reference to the others.

Platform reporting systems — Meta Ads Manager, Google Analytics, TikTok for Business — report performance as measured by the platform itself, against the spend processed through that platform's billing system. Agency management fees, creative production costs, and tool subscriptions that enable the platform spend are invisible to these systems. The reported return on ad spend is calculated against a fraction of the true investment.

Departmental reporting compounds the fragmentation. Each function produces performance reports in its own format, using its own metrics, selected and framed by the department being assessed. The cross-departmental, independent view of whether total organisational expenditure is producing proportionate outcomes is structurally absent from standard reporting frameworks.

1.3 Information Overload and Underutilisation

The consequence of this fragmentation is not a shortage of data but an excess of it — presented in forms that obscure rather than illuminate the relationship between cost and outcome. Business leaders routinely receive marketing dashboards showing impressions, reach, click-through rates, and attributed conversions; HR reports showing training completion rates and time-to-hire; logistics reports showing fulfilment speed and carrier performance — while remaining unable to answer whether the aggregate investment across all these functions is producing the business results that justify it.

This phenomenon — data abundance coexisting with insight scarcity — has been documented across sectors. The proliferation of business intelligence tools, dashboard platforms, and analytical infrastructure has not resolved the problem. It has in many cases amplified it by creating additional reporting layers that require interpretation without providing the independent, executive-level synthesis that business leaders require for capital allocation decisions.

1.4 The Missing Function

What is structurally absent from the contemporary business reporting landscape is an independent, cross-departmental, statistically grounded measurement of the relationship between what organisations spend and what they produce. Not platform reporting — which measures performance within a platform's own ecosystem. Not departmental activity reporting — which measures what a function did rather than what the investment in it produced. Not financial recording — which captures expenditure without connecting it to outcomes.

The cost-outcome intelligence layer that sits above all of these and connects them to executive decision-making is the gap that Cost Impact Analysis is designed to address. Its absence is not a technical problem — the statistical methods required to produce it are well established. It is an accessibility and implementation problem: the existing methods are either inaccessible to most organisations at reasonable cost, or oriented toward specific domains rather than the cross-departmental executive view that capital allocation decisions require.

2. Evidence Across Sectors

The cost-outcome measurement gap is not a characteristic of any single sector but a structural feature of how organisations of all types report and interpret their expenditure. The specific manifestations differ across private, public, and civil society contexts — but the underlying gap is consistent across all three.

2.1 Private Sector

In commercial organisations the cost-outcome measurement gap manifests most acutely in marketing and digital operations, where the conflict between platform-reported performance and true business return is most directly observable. A retail business with 500,000 euros in annual digital marketing spend may receive platform reports indicating a 4x return on ad spend — a figure that excludes the agency management fee, the monthly creative production retainer, the tool subscription stack, and the freelancer costs that together represent the true loaded cost of the digital operation. When all costs are included, the effective return against total investment may be materially different from the platform-reported figure.

The same gap appears across other functional areas. HR investment in training and development is routinely approved without a defined framework for connecting it to productivity, retention, or performance outcomes. Logistics cost growth in e-commerce operations frequently outpaces the fulfilment performance improvements it is supposed to enable. Operational expenditure — facility costs, service agreements, software subscriptions — accumulates through automatic renewal without systematic review of whether the cost is producing proportionate output.

The private sector case is strengthened by the documented scale of the problem. When 44 percent of SaaS licenses go unused and digital advertising costs rise at double the rate of general inflation, the cost of measurement failure is not theoretical — it is the quantifiable difference between what organisations spend and what they could spend if allocation decisions were evidence-based.

2.2 Public Sector

Public sector organisations operate under explicit accountability obligations that make the cost-outcome measurement gap particularly consequential. Public resources require demonstrable outcome connection — not as an optional governance improvement but as a fundamental requirement of public financial management.

Yet the analytical frameworks available to most public sector finance functions are structurally oriented toward expenditure recording and budget compliance rather than outcome measurement. Programme costs are tracked against budget lines. Service delivery is reported in activity terms — number of services provided, cases processed, participants reached. The systematic statistical connection between what a programme costs and what it produces in terms of citizen outcomes, service quality, or operational efficiency is rarely examined continuously or with rigorous methodology.

The growing pressure from audit bodies, central government, and public accountability frameworks for value-for-money evidence exists in an environment where the analytical infrastructure to produce that evidence is underdeveloped. Cost Impact Analysis addresses this gap directly — providing a methodology that connects programme expenditure to defined outcome metrics within the analytical capability constraints of lean public sector finance functions.

2.3 Civil Society

Non-profit and non-governmental organisations face a version of this problem that is uniquely complicated by dual accountability — to donors for how funds were spent, and to beneficiaries and mission for what those funds produced. Donor reporting frameworks typically require evidence of expenditure. What they less consistently require, and what organisations less consistently produce, is systematic statistical evidence of what that expenditure produced in terms of programme outcomes.

The tension between administrative cost minimisation — itself a legitimate donor expectation — and the investment in analytical capability required to produce rigorous outcome evidence creates a structural constraint. Organisations with lean finance functions and limited analytical resource cannot implement enterprise-grade measurement frameworks. Cost Impact Analysis provides a methodology that operates within these constraints — requiring consistent data input rather than sophisticated infrastructure, and producing statistically grounded output accessible to non-specialist users.

3. Critique of Existing Approaches

Current solutions to the cost-outcome measurement problem are either inaccessible, insufficient, or structurally conflicted. The response of increasing analytical complexity is in many cases counterproductive rather than curative. This section examines each major existing approach on its own terms before identifying the specific gap that none of them address.

3.1 Platform and Channel Reporting

The most widely used source of cost-outcome information in digital operations is platform reporting — the performance data provided by advertising and marketing platforms about the campaigns run through their systems. This reporting has a structural conflict of interest that fundamentally limits its utility as an independent cost-outcome measurement tool.

Platforms have a commercial incentive to report metrics that justify continued and increased advertising spend. Attribution models that credit the platform's own inventory, reporting windows calibrated to maximise attributed conversion counts, and the systematic exclusion of surrounding costs from ROI calculations all serve this incentive. The result is a measurement framework designed by the entity being paid to deliver results, using metrics calibrated to demonstrate the value of that payment. Evidence from a US antitrust trial confirmed that Google Ads pricing has consistently increased over time, with internal documents suggesting awareness that price increases would be absorbed regardless of measured performance (PPC Hero, 2025). Its limitations as an independent cost-outcome measurement tool are structural rather than incidental.

3.2 Business Intelligence Infrastructure

Business intelligence platforms — including Power BI, Tableau, Looker, and enterprise equivalents — offer genuine analytical power for organisations that can deploy them effectively. Their limitations in the context of accessible cost-outcome measurement are practical rather than technical.

Effective BI deployment requires data engineering to connect and normalise data sources, analytical expertise to build and maintain models, and ongoing technical resource to update and interpret outputs as business conditions change. These requirements are manageable for large organisations with dedicated data functions. For the majority of SMBs, public sector bodies with lean finance functions, and civil society organisations operating under administrative cost constraints, the resource requirement to implement and sustain effective

BI infrastructure exceeds what is realistically available. The result is that the organisations most in need of independent cost-outcome intelligence are the least able to access it through existing BI frameworks.

3.3 Marketing Mix Modelling

Marketing Mix Modelling represents the most statistically rigorous approach currently available for connecting marketing expenditure to business outcomes. Its methodology — regression-based analysis of the relationship between marketing inputs and revenue or other outcome variables — is conceptually aligned with the Cost Impact Analysis approach proposed in this paper. Its limitations are primarily those of accessibility and scope.

MMM implementations at the enterprise level require significant data science expertise, substantial historical data, and analytical resource that places them beyond the reach of most SMBs and non-commercial organisations. The open-source frameworks released by Meta (Robyn) and Google (Meridian) have democratised the statistical methodology without democratising its implementation — the technical barrier to effective deployment remains substantial for organisations without dedicated analytical resource.

Additionally, MMM is designed primarily around media and marketing spend, with a focus on channel attribution and budget optimisation. Its application across non-marketing cost categories — HR investment, operational expenditure, logistics costs, public sector programme spend — is limited by both its design orientation and the data models on which it depends.

3.4 Management Accounting Frameworks

Activity-based costing, balanced scorecard, and related management accounting approaches provide valuable frameworks for cost allocation and strategic performance measurement. Their contribution to organisational financial management is well established in the academic and practitioner literature (Kaplan and Norton, 1992; Cooper and Kaplan, 1988).

Their limitation in the specific context of continuous, accessible cost-outcome measurement is that they are typically implemented as periodic strategic exercises rather than continuous operational frameworks, require specialist accounting expertise to design and maintain, and focus on cost allocation and strategic alignment rather than the statistical relationship between cost changes and outcome changes over time. They address the question of how costs should be allocated rather than the question of whether the costs being incurred are producing the outcomes they are supposed to support.

3.5 The Complexity Trap

A consistent pattern across organisational responses to the cost-outcome measurement gap is the tendency to respond to data fragmentation with increased analytical complexity — adding reporting layers, dashboard tools, and measurement frameworks that compound rather than resolve the underlying problem. Each additional tool generates additional data requiring additional interpretation. The organisations most overwhelmed by reporting fragmentation frequently have the most sophisticated reporting infrastructure.

This pattern suggests that complexity is not the appropriate solution to the problem of measurement fragmentation. The analytical framework that best serves executive decision-making is not the most comprehensive or technically sophisticated — it is the one that most directly connects the question being asked to a statistically grounded answer, with the minimum technical complexity required to produce reliable output. This principle underlies the design of Cost Impact Analysis.

4. The Proposed Methodology — Cost Impact Analysis

4.1 The Business Optimisation Context

Cost Impact Analysis is proposed as a business optimisation discipline — a systematic methodology for connecting organisational expenditure to measurable business outcomes through regression-based correlation analysis. It operates at the intersection of management accounting, applied statistics, and business intelligence, serving the cost-outcome measurement need that existing approaches in each of these fields address only partially.

As a business optimisation discipline, Cost Impact Analysis is distinguished from marketing analytics, campaign performance measurement, and financial recording by its orientation toward executive decision-making rather than operational reporting. Its primary output is not an account of what a function did or what a platform delivered — it is a statistically grounded view of the relationship between what an organisation invested and what it produced, independent of the reporting frameworks of the departments and vendors involved.

4.2 The Core Proposition

The core proposition of Cost Impact Analysis is that bypassing the micro-metric layer — the impressions, clicks, completion rates, satisfaction scores, and activity indicators that constitute the majority of contemporary organisational reporting — and observing the direct statistical relationship between categorised costs and defined outcome KPIs produces more useful executive intelligence than any accumulation of granular activity data.

This is not an argument against granular operational metrics. Those metrics serve legitimate operational management purposes within their respective functions. It is an argument that the question relevant to executive decision-making — whether the aggregate investment in a function is producing proportionate business outcomes — is best answered at the level of the cost-outcome relationship rather than through the accumulation of activity indicators that each function selects to represent its own performance.

4.3 Analytical Architecture — A Two-Layer Statistical Framework

Cost Impact Analysis employs a deliberate two-layer statistical architecture in which different analytical methods are applied to different analytical questions. The primary layer uses multivariate ordinary least squares regression to estimate the independent contribution of each cost category to KPI outcomes. The secondary layer uses Pearson correlation on

transformed series to assess co-movement strength for visual representation. The two layers are intentionally distinct in their pre-processing, their mathematical form, and their interpretive purpose. This section documents each layer precisely.

4.3.1 Primary Layer — Multivariate Ordinary Least Squares Regression

The primary analytical engine of Cost Impact Analysis is multivariate ordinary least squares regression, in which the selected outcome KPI is modelled as a linear function of all cost category expenditures observed in the same period.

For a given observation window containing t aligned periods, the model takes the form:

$$Y_t = \beta_0 + \beta_1 X_{1t} + \beta_2 X_{2t} + \cdots + \beta_k X_{kt} + \epsilon_t$$

Where Y_t is the KPI value in period t, X_{1t} through X_{kt} are the total expenditures per cost category k in period t, β_0 is the intercept, β_1 through β_k are the regression coefficients expressing KPI movement per unit of currency spent in each category, and ϵ_t is the residual error term.

The model is estimated by minimising the sum of squared residuals across all aligned periods. The multivariate formulation is the preferred specification because it estimates the partial effect of each cost category on the KPI while holding all other categories constant. This directly addresses the confounding variable problem inherent in business expenditure data, where multiple cost categories frequently co-move — for example, during a seasonal campaign push where platform spend, agency hours, and production costs increase simultaneously.

Pairwise analysis of any single category against the KPI in such circumstances would conflate the individual category effect with the shared movement of all co-varying categories. The multivariate OLS model isolates each coefficient estimate from this confounding. The pipeline from raw data to regression output is: aggregate expenditure per category per period from invoice-level records; match each period to its corresponding KPI observation; retain only periods for which a KPI value exists; submit the aligned panel to OLS estimation with intercept. No detrending or differencing is applied to the series prior to regression.

The regression operates on raw period-level values, which preserves the interpretability of the slope coefficients as KPI movement per unit of currency spent in that category during that period. An optional spend attribution mode — timespan distribution — spreads each invoice payment across the days until the subsequent payment in the same category, rather than attributing the full invoice amount to the period of the invoice date.

This mode changes how expenditure is allocated across periods and may better reflect the operational consumption pattern of retainer-based or subscription costs. It is a data preparation choice rather than a statistical transformation: the regression operates on raw aligned levels in either mode.

The primary outputs of the regression layer are the unstandardized OLS coefficient β_k for each cost category, expressing the estimated KPI movement per unit of currency spent holding all other categories constant; the coefficient of determination R^2 for the overall model, expressing the proportion of total KPI variance explained by the full set of cost categories; and a standalone bivariate Pearson correlation coefficient r computed independently for each category to capture unadjusted, isolated historical co-movement.

To facilitate cross-category comparison regardless of differing spend scales, the system computes the standardized regression coefficient β_k^* for each category:

$$\beta_k^* = \beta_k \times \left(\frac{\sigma_{X_k}}{\sigma_Y} \right)$$

The methodology alters its ranking, sorting, and impact-labeling mechanics based on the active analytical mode: Multivariate Mode: Categories are sorted and prioritized by the absolute value of their standardized regression coefficient (β_k^*), which isolates independent impact while controlling for category variability.

The bivariate correlation r is retained strictly as a secondary visual indicator to show standalone, unadjusted historical co-movement. Per-Category Fallback Mode: Where data constraints trigger the univariate fallback, sorting, ranking, and impact labeling default strictly to the magnitude of the standalone bivariate Pearson correlation coefficient r.

Where the number of cost categories exceeds the available degrees of freedom, or where multicollinearity among categories is too severe for stable multivariate estimation, the

system falls back to univariate regression for each category independently. In the univariate specification, each category's spend series is regressed against the KPI series in isolation, producing a separate slope estimate and correlation coefficient for each category without controlling for co-movement between categories. Results from the univariate fallback should be interpreted with corresponding caution regarding confounding.

4.3.2 Secondary Layer — Cost Flow Correlation on Transformed Series

The secondary analytical layer serves a distinct purpose: visual representation of co-movement strength between cost categories and the KPI after removing long-term trends and period-on-period persistence. This layer powers the Cost Flow Analysis visualisation and uses a stricter transformation pipeline before applying Pearson correlation. Each cost series and the KPI series are independently subjected to two pre-processing steps before correlation is computed.

First, linear detrending removes the long-term trend component from each series using ordinary least squares regression of the series against a time index:

$$detrended(t) = value(t) - (\alpha + \beta t)$$

Where α and β are estimated by OLS. The detrended series represents the deviation of each observation from its predicted trend value.

Applied independently to both the cost series and the KPI series, detrending ensures that subsequent correlation reflects genuine co-movement rather than the coincidence of two variables that both happen to increase over time.

Second, first differencing converts each detrended series from levels to period-on-period changes:

$$diff(t) = detrended(t) - detrended(t-1)$$

First differencing further reduces autocorrelation in the series and focuses the correlation analysis on short-term co-movement between cost changes and KPI changes rather than

level relationships. One observation is lost per series through differencing; accordingly, the minimum data requirement for this layer is nine raw aligned periods, yielding eight usable differenced observations.

Pearson correlation is then computed on the differenced, detrended series. For variables X (cost) and Y (KPI), after the transformation pipeline, the Pearson coefficient r is:

$$r = \frac{\sum(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum(X_i - \bar{X})^2 \sum(Y_i - \bar{Y})^2}}$$

Where $\bar{X}$ and $\bar{Y}$ denote the sample means of the transformed series respectively. Statistical significance is assessed using the standard t-transformation for Pearson r with $n - 2$ degrees of freedom:

$$t = \frac{r\sqrt{n-2}}{\sqrt{1-r^2}}$$

The associated two-tailed p-value is computed under the null hypothesis of no linear relationship.

Only correlations meeting a significance threshold of $p \leq 0.05$ are rendered as connections in the Cost Flow visualisation; below-threshold relationships are suppressed or faded, ensuring that the visual output reflects statistically defensible relationships rather than noise.

The intentional distinction between the two layers is methodologically significant. The regression layer (Section 4.3.1) operates on raw levels and estimates interpretable slope coefficients suited to the question “how much does the KPI move per unit of currency spent in this category?”

The correlation layer operates on detrended differenced series and addresses the stricter question “do spend and KPI co-move after removing shared trends and period persistence?” The two questions are related but not equivalent. Regression on levels can be influenced by shared long-term growth trends; correlation on differenced series specifically removes that influence. Each layer is the appropriate tool for its designated purpose.

4.4 Application Requirements

Cost Impact Analysis produces reliable output under specific conditions that define its appropriate scope of application.

Variable costs with clear category definition — costs that change meaningfully period to period and are consistently attributed to a defined category — provide the variation necessary for both the regression and correlation layers to identify relationships. Fixed costs with no period-on-period variation, such as administrative salaries that remain constant across the observation period, produce no regression coefficient signal and no correlation output, and are outside the productive scope of the methodology.

Defined, measurable outcome KPIs — business metrics recorded consistently at the same frequency as the cost data — are required on the outcome side of the analysis. The methodology requires a quantified outcome variable against which cost variation can be examined. Outcomes that are not quantified or not recorded consistently are outside the methodology's scope.

Sufficient data history is required for both analytical layers, with differing minima reflecting their different statistical structures. The regression layer requires a minimum of eight weekly periods or six monthly periods with aligned KPI observations, ensuring sufficient degrees of freedom for stable OLS coefficient estimation. The Cost Flow correlation layer requires a minimum of nine raw aligned periods, producing eight usable observations after first differencing. Results become progressively more reliable with longer observation periods; twelve or more months of consistent data provides meaningfully more stable outputs from both layers.

4.5 Constraints and Limitations

The methodology's limitations are as important to document as its capabilities. Cost Impact Analysis identifies statistical relationships between cost changes and KPI changes — it does not establish causal direction or isolate the effect of a single cost category from other variables that influence the outcome simultaneously.

Statistical association is not causation. A strong regression coefficient or correlation between a cost category and a KPI may reflect a genuine cost-outcome relationship, shared influence from an unobserved variable such as seasonal patterns affecting both costs and outcomes simultaneously, or coincidental co-movement over the observation period. The multivariate regression layer reduces but does not eliminate this risk by controlling for co-varying cost

categories. Findings from both layers should be treated as signals warranting further investigation rather than conclusions warranting immediate action.

Highly seasonal businesses with short operating periods require multiple complete cycles of data before seasonal patterns can be distinguished from the underlying cost-outcome relationship. Cost categories that are highly correlated with each other — for example, platform spend and agency management fees that always move in proportion — may produce similar regression coefficients and correlation outputs. While the multivariate OLS specification controls for this co-movement in coefficient estimation, severe multicollinearity among categories reduces the precision of individual coefficient estimates and may trigger the univariate fallback specification.

4.6 The Spend Efficiency Construct

Cost Impact Analysis incorporates a composite metric — the Spend Efficiency Score — that synthesises multiple dimensions of cost-outcome performance into a single directional indicator expressed on a normalised scale of zero to one hundred.

The construct combines three components: the ratio of KPI growth rate to cost growth rate over the observation period; the aggregate correlation strength between cost categories and outcome KPIs; and the concentration of spend across categories weighted by their correlation direction. Each component is normalised before combination to ensure comparability across different scales of measurement.

The composite score provides an executive-level indicator of whether cost-outcome efficiency is improving or declining over time — a directional signal that complements the more granular category-level analysis without requiring the reader to synthesise multiple data points simultaneously.

5. Implications for Business Efficiency & Optimisation

5.1 The Executive Intelligence Layer

Cost Impact Analysis addresses a structural gap in the information available to executive decision-makers — providing independent, cross-departmental cost-outcome intelligence that departmental reporting, platform measurement, and financial recording each address only partially and in isolation.

The methodology is positioned explicitly on the CEO-CFO axis — serving the financial oversight and resource allocation needs of senior leadership rather than the operational reporting needs of individual departments. Its independence from departmental reporting frameworks is both a design characteristic and a functional requirement: an assessment of whether a department's investment is producing adequate returns cannot be produced reliably by the department being assessed.

5.2 Sector-Specific Implications

In commercial organisations Cost Impact Analysis addresses the capital allocation efficiency question — whether the total investment across all monitored functions is producing proportionate business returns — at an accessible price point and with a technical requirement commensurate with what SMBs and mid-size enterprises can realistically sustain.

In public sector organisations it addresses the value-for-money evidence requirement that audit frameworks and accountability obligations increasingly demand, providing a continuous monitoring capability rather than a periodic audit intervention.

In civil society organisations it addresses the dual accountability challenge — providing mission outcome evidence to donors and funders while operating within the administrative cost constraints that characterise lean non-profit finance functions.

5.3 The Democratisation of Analytical Capability

A recurring theme in the critique of existing approaches is the gap between the statistical rigour available to large organisations with dedicated analytical resource and the measurement tools accessible to smaller organisations without it. Marketing Mix Modelling at the enterprise level, BI infrastructure requiring data engineering, and management

accounting frameworks requiring specialist expertise all occupy the sophisticated end of this spectrum.

Cost Impact Analysis is designed explicitly for the accessible end — requiring consistent data input rather than complex infrastructure, producing outputs interpretable by financial non-specialists, and calibrated for deployment across organisations that cannot justify enterprise-grade analytical investment. The democratisation of rigorous cost-outcome analysis is both a design objective of the methodology and, at aggregate scale, a contribution to economic productivity by reducing the analytical capability gap between large and small organisations.

5.4 Contribution to Business Optimisation Literature

This paper makes the following contributions to the academic literature at the intersection of business optimisation, management accounting, applied statistics, and business intelligence:

It proposes Cost Impact Analysis as a named, documented methodology within the business optimisation literature — addressing a gap that existing frameworks in management accounting (Journal of Management Accounting Research, Management Accounting Research), business intelligence (Decision Support Systems, Journal of Business Analytics), and marketing accountability (Journal of Marketing, Journal of the Academy of Marketing Science) each address partially but none address completely in the accessible, cross-departmental form proposed here.

It identifies and documents the cost-outcome measurement gap as a structural feature of contemporary business reporting rather than a sector-specific or organisation-specific phenomenon — present across private, public, and civil society contexts in different manifestations but rooted in the same structural disconnection between expenditure recording and outcome measurement.

It proposes a two-layer analytical architecture — multivariate OLS regression on raw aligned period levels as the primary coefficient estimation engine, complemented by Pearson correlation on detrended first-differenced series for visual co-movement representation — as an accessible, statistically sound methodology for cost-outcome analysis, distinct from MMM in its cross-departmental scope, accessibility to non-specialist users, and separation of interpretable slope estimation from trend-adjusted connection strength visualisation.

It introduces the Spend Efficiency Score as a composite construct for synthesising multiple dimensions of cost-outcome performance into a single directional executive metric — a

contribution to the business optimisation measurement literature that addresses the synthesis problem accompanying multi-variable analytical frameworks.

Most significantly, it frames the cost-outcome measurement gap as an economic efficiency problem rather than solely a management methodology gap — connecting the aggregate consequences of measurement failure to documented literature on capital misallocation, platform concentration, and the SMB productivity gap, and positioning the proposed methodology as a contribution to both fields simultaneously.

Conclusion

The gap between what organisations spend and what they can demonstrate that spending produced is a structural feature of contemporary business reporting — not a failure of individual organisations or sectors, but a consequence of reporting frameworks designed around activity measurement rather than outcome measurement, cost recording rather than cost-outcome connection, and departmental performance rather than cross-departmental financial intelligence.

At aggregate scale this gap constitutes a market inefficiency with measurable macroeconomic consequences — directing capital toward self-reported activities rather than productive outcomes, enabling systematic value transfer to platform intermediaries, widening the productivity gap between analytically capable and analytically constrained organisations, and limiting the ability of public institutions to demonstrate value for money against growing accountability obligations.

Existing approaches to this problem — platform reporting, business intelligence infrastructure, marketing mix modelling, and management accounting frameworks — each address aspects of it within specific contexts and at accessibility levels that exclude the majority of organisations most in need of independent cost-outcome intelligence.

Cost Impact Analysis is proposed as a business optimisation methodology that addresses the gap directly — connecting categorised organisational expenditure to defined outcome KPIs through regression-based correlation analysis, with a technical requirement and cost structure accessible to organisations across private, public, and civil society sectors.

The methodology does not claim to replace more sophisticated analytical approaches where those are available and appropriate. It claims to provide a rigorous, accessible, independent cost-outcome intelligence layer for the significant proportion of organisations for which existing approaches are not a realistic option — and to do so within a documented, statistically grounded framework that meets the evidence standards required by the executive, governance, and accountability contexts in which it is designed to operate.

Further research, including empirical validation of the methodology across multiple organisational contexts and sectors, comparative analysis of its outputs against alternative measurement approaches, and longitudinal studies of the relationship between Spend Efficiency Score trends and organisational performance outcomes, would strengthen the evidence base for Cost Impact Analysis as a standard business optimisation discipline.

References

Statistical Methodology

Box, G.E.P. and Jenkins, G.M. (1976) Time Series Analysis: Forecasting and Control. San Francisco: Holden-Day.

Granger, C.W.J. (1969) "Investigating Causal Relations by Econometric Models and Cross-spectral Methods." Econometrica, 37(3), pp. 424–438.

Kendall, M. and Stuart, A. (1979) The Advanced Theory of Statistics. London: Griffin.

Pearson, K. (1895) "Notes on Regression and Inheritance in the Case of Two Parents." Proceedings of the Royal Society of London, 58, pp. 240–242.

Business Intelligence and Analytics

Davenport, T.H. and Harris, J.G. (2007) Competing on Analytics: The New Science of Winning. Boston: Harvard Business School Press.

Elbashir, M.Z., Collier, P.A. and Davern, M.J. (2008) "Measuring the Effects of Business Intelligence Systems." International Journal of Accounting Information Systems, 9(3), pp. 135–153.

Negash, S. (2004) "Business Intelligence." Communications of the Association for Information Systems, 13(1), pp. 177–195.

Management Accounting

Cooper, R. and Kaplan, R.S. (1988) "Measure Costs Right: Make the Right Decisions." Harvard Business Review, 66(5), pp. 96–103.

Kaplan, R.S. and Norton, D.P. (1992) "The Balanced Scorecard — Measures that Drive Performance." Harvard Business Review, 70(1), pp. 71–79.

Otley, D. (1999) "Performance Management: A Framework for Management Control Systems Research." Management Accounting Research, 10(4), pp. 363–382.

Marketing Mix Modelling

Borden, N.H. (1964) "The Concept of the Marketing Mix." Journal of Advertising Research, 4(2), pp. 2–7.

Jin, Y., Wang, Y., Sun, Y., Chan, D. and Koehler, J. (2017) "Bayesian Methods for Media Mix Modeling with Carryover and Shape Effects." Google Research Technical Report.

Leeflang, P.S.H., Verhoef, P.C., Dahlstrom, P. and Freundt, T. (2015) "Challenges and Solutions for Marketing in a Digital Era." European Management Journal, 32(1), pp. 1–12.

Economics and Capital Allocation

Hsieh, C.T. and Klenow, P.J. (2009) "Misallocation and Manufacturing TFP in China and India." Quarterly Journal of Economics, 124(4), pp. 1403–1448.

OECD (2021) OECD SME and Entrepreneurship Outlook 2021. Paris: OECD Publishing.

Philippon, T. (2019) The Great Reversal: How America Gave Up on Free Markets. Cambridge: Belknap Press.

Syverson, C. (2011) "What Determines Productivity?" Journal of Economic Literature, 49(2), pp. 326–365.

Industry Data

BetterCloud (2024) State of SaaSOps 2024. Available at: bettercloud.com.

Colorlib (2026) "130+ SaaS Statistics: Market Size, Growth & Trends." Available at: colorlib.com.

Digital Silk (2025) "47 SaaS Statistics Every Business Leader Should Know." Available at: digitalsilk.com.

PPC Hero (2025) "Why are Google Ads CPCs Increasing?" Available at: ppchero.com.

SellersCommerce (2026) "27 Eye-Opening SaaS Statistics in 2025." Available at: sellerscommerce.com.

Triple Whale (2025) DTC Ad Benchmark Report 2025. Available at: triplewhale.com.

TruLata Solutions (2025) "Navigating Google CPC Inflation in 2025." Available at: trulata.com.

WordStream (2024) Google Ads Industry Benchmarks 2024. Available at: wordstream.com.